(72) Inventors:
• **Röhrle, Klaus, Sony International Europe GmbH 70736 Fellbach (DE)**
• **Kovacs, Ernö, Sony International Europe GmbH 70736 Fellbach (DE)**

(74) Representative: **Rupp, Christian, Dipl.Phys. et al Mitscherlich & Partner Patent- und Rechtsanwälte Sonnenstrasse 33 80331 München (DE)**

(54) **Management of a communication network and the migragtion of mobile agents**

(57)   A technique for managing the migration of mobile agents to nodes of a communication network is proposed. The trustworthiness of at least one node (102b) of the network is checked (301). In case the trustworthiness exceeds a pre-set trust threshold, a trust token for the checked node (102b) is generated (306) and the trust token is stored (303) in the network. In advance to a migration of a mobile agent (104a) to a node of the network it is verified (108, 109) if a valid trust token is existing for the corresponding node (102b). The migration of the mobile agent (104a) is restricted (107) to nodes having a valid trust token.

Fig 4.

BEST AVAILABLE COPY

EP 1 067 457 A1

Description

[0001]   The present invention relates to a communication network, a method for managing a communication network and method for managing the migration of mobile agents to nodes of a communication network.

[0002]   The present invention has its application in the field of communication networks, and preferably in the field of mobile computing, wireless communication, mobile multimedia middleware and so-called mobile agents. One particular application is the control (management) of the migration of so-called mobile agents between places in the network.

[0003]   Communication is concerned with providing networked services over (for example wireless) networks to the users. So-called mobile agents is a new computing concept enabling the transfer of an executing piece of software (code) from one node of the computer network to another node. One example for a mobile code system are Java applets loaded over the Internet in a Web browser system.

[0004]   The concept of mobile agents is set forth in detail in US-A-5 603 031. Regarding the general concept of the mobile agent technology therefore reference is made to said prior art.

[0005]   Figure 1 shows the basic operation of a mobile agent system. An agent system offers mobile agents an execution environment which is called a place. Mobile agents are a piece of executing software being capable of migrating from one agent system to another. On all places mobile agents can access local services offered.

[0006]   Figure 1 shows the migration of a mobile agent 104a between two agent systems 101a, 101b. The agent systems contain execution environments which are called places 102a, 102b. In the places local services 103a, 103b are provided offering services 106 to the mobile agent 104a. In the migration process 105 the mobile agent is serialised on place 102a and the state and the code of the agent is transferred to the place 102b. In the place 102b the mobile agent 104a is recreated from the transferred state and code. Then the execution thread of the mobile agent 104a can be executed in the place 102b.

[0007]   In the current state of the art agents are migrating between computing nodes on the Internet. There are several security-related problems in conjunction with mobile agents, e.g. to protect the agent place against a malicious agent. One hard problem is the protection of an agent against a malicious host.

[0008]   A malicious host can perform several kinds of attacks to a mobile agent like e.g. spy out its code or data, manipulate its code or data, execute it incorrectly or deny the execution at all. At the moment there are no known security mechanisms, which protect an agent against all these attacks except of using trusted hardware.

[0009]   A malicious host is only a problem for agents whose owners cannot trust the place operators. Trust means in this context that the owner of the agent knows (or at least hopes) that the place will not attack one of its agents. As there is no solution for protecting agents against malicious places, there are some approaches to circumvent this problem by allowing agents only to move to hosts trusted by the owner of the mobile agent.

[0010]   These approaches are especially important when the itinerary of an agent is not known in advance. In most agent environments an agent can retrieve during runtime the location of certain services and then visit these locations. In this case the agents itinerary is determined dynamically. This is in contrast to a static approach to inform a mobile agent in advance concerning which places are not rusted by the owner of the mobile agent. It is to be noted that even the trustworthiness places can change in time, i.e. a trustworthy place can become non-trustworthy later on which results in problems when informing a mobile agent in advance about the trustworthiness of places to be vissited.

[0011]   A mobile agent owner might want to restrict the agent in a way that it is not allowed to migrate to places the mobile agent owner does not trust. To achieve this, the agent needs a means to find out if a given place is trusted by his user or not.

[0012]   This can be done either by an ,,organisational" approach, which means that only trustworthy parties can start a place. The disadvantage of this approach is that it leads to an agent environment, which is not ,, open" anymore (open in the sense that everybody can start a place).

[0013]   Another approach is to specify the trusted places an agent is allowed to migrate in advance. In this case, the agent carries a list of trusted places and either the agent system or the agent itself checks on a migration whether the agent is allowed to go to its destination or not.

[0014]   This approach also has a big disadvantage: it is not always clear in advance whether a place is trusted or not, which can severely reduce the number of places an agent is allowed to migrate. A user normally 'knows' only a small number of places in the agent environment. He can of course only judge agent systems he knows. Therefore he will probably not trust most of the places which means that his agents are only allowed to migrate to a very limited set of places. This may limit the usefulness of an agent environment significantly.

[0015]   In view of the above prior art it is the object of the present invention to provide for a technique allowing for a facilitated trust check in networks. It is particularly the object of the invention to provide such a trust check for the management of the migration of mobile agents in a communication network.

[0016]   The trust check technique should furthermore be open for a dynamic in-line trust check approach.

[0017]   The above object is achieved by means of the features of the independent claims. The dependent claims develop further the central idea of the invention.

[0018]   According to the present invention therefore a

method for managing a communication network is provided. The trustworthiness of at least one node of the network is checked. The trustworthiness is compared to a pre-set threshold and in case the trustworthiness exceeds the pre-set trust threshold, a trust token for the checked node is generated and the trust token indicating the fact that the corresponding node has passed the trustworthiness check is stored in the network.

[0019] The validity of a token can be ensured by cryptographic measures such as f.e. difital signatures.

[0020] The step of checking the trustworthiness of at least one node of the network can be executed by another node of the network having a trust centre function which is called trust centre in the following description.

[0021] A predetermined expiry date limiting the life span of a trust token can be added to the trust token.

[0022] According to another aspect of the invention a method for managing the migration of mobile agents to nodes of a communication network is provided. The trustworthiness of at least one node of the network is checked. In case the trustworthiness exceeds a pre-set trust threshold, a trust token for the checked node is generated and stored in the network. In advance to migration of mobile agent to a node of the network, it is verified if a valid trust token is existing for the corresponding node. The migration of the mobile agent is restricted to nodes having a valid trust token.

[0023] The verification of the trust token can be executed dynamically during the run time of a mobile agent.

[0024] The step of checking the trustworthiness of at least one node of the network can be executed by at least one third node of the network having a trust centre function.

[0025] A list of names of trust centre nodes can be attached to a mobile agent. The list can be protected against attempts of manipulation by cryptographic measures such as f.e. digital signatures.

[0026] The step of verification if a valid trust token is existing for a corresponding node in advance to a migration of the mobile agent can be effected by querying the corresponding target node for valid trust token. Additionally or alternatively at least one trust centre node can be queried and/or at least one storage unit storing valid trust token and the corresponding nodes of the network can be queried.

[0027] According to the present invention furthermore a communication network is provided. At least one trust centre node of a network checks the trustworthiness of at least one node of the network. The trust centre node has a token creator for generating a trust token for the checked node in case the trustworthiness exceeds a pre-set trust threshold. The trust centre can be another node of the network as each checked node. The trust token and the corresponding checked nodes can be stored in a trust manager database being part of or being connected to the trust centre node.

[0028] At least one mobile agent place can be provided in the network for managing the migration of mobile

agents in the communication network. The mobile agent places can be signed to access the trust manager database of at least one trust centre node.

[0029] The mobile agent place can comprise a local database to store information fetched from the trust manager database of a trust centre node.

[0030] A list of trust centre nodes can be attached to a mobile agent managed by a mobile agent place.

[0031] In the following preferred embodiments of the present invention will be explained with reference to the figures of the enclosed drawings.

Figure 1 shows the migration of a mobile agent,

Figure 2 shows a wireless computing scenario,

Figure 3 shows the relationship between a trust centre node, an agent place and a directory service,

Figure 4 shows the system architecture for a trust centre,

Figure 5 shows a trust token with some related data fields, and

Figure 6 explains the architecture of a agent system.

[0032] Figure 1 has already been explained in the introductory part of the descritpion.

[0033] Figure 2 shows a typical wireless computing scenario. Several mobile devices are connected over a wireless network link to a base station. They use different wireless network equipment. The base stations are connected over a network (e.g. the Internet) with each other and with network servers.

[0034] Figure 3 shows the relationship between the trust centre a agent place and a directory service. The trust centre generates a trust token and issues it to the agent place. This place can then declare its trustworthiness using the issued trust token. Agent wishing to migrate can request the trust tokens from several places (the target agent place, the trust centre, a directory or others). The checks can be performed by the agent or by the agent system.

[0035] A place may get trust tokens from various trust centres. Trust token may also be made public by the trust centres or the place, so everybody can verify that a place is trusted by a specific trust centre ("Publishing of the trust centre"). This can be done through the use of a directory service like the international X.500 directory. Figure 3 explains this process.

[0036] Figure 4 shows the system architecture for a trust centre. The trust centre contains a token creator module, which creates a trust token after the credibility of a agent place is checked. It used cryptographic and digital signatures for ensuring the validity of his statement. The trust token is managed by the trust manager

module which stores the trust token and addition information in a trust token database. The revoke manager module is a supplementary module which enables to define tokens as revoked after they are issued. Although this is not a 100% solution it increases the likelihood of preventing accidental migration to untrusted places. The revoke list can be distributed to interested agent places to improve the checking of the trust tokens. The token requester module handles request for issued trust tokens from agent places. The directory access module allows access to a directory service.

[0037] There is a plurality of implementation possibilities for the revoking procedure:

- a mobile agent is provided with a revoke list
- the trust centre is accessed to query revoked places, and/or
- a central revoke manager is provided

[0038] As shown in figure 4 a trust centre consists of several modules, each performing a specific task. The token creator module creates the trust token. The trust token consists of some information about the trusted agent place. This information is digital signed with the private key of the trust centre so that user of the trust token can verify that the content of it was created by the trust centre.

[0039] On the side of the agent place, trust tokens are managed by the trust token manager. This module can keep a cache of already available trust tokens. The revoke list manager module stores a list of revoked trust tokens. It can synchronise its content with the revoke list database from the trust centre.

[0040] The information, if a place is trustworthy or not, is determined dynamically during the runtime of the agent. This is achieved by providing adding trusted third parties - called trust centres. An owner of an agent cannot only specify places he trusts but also trust centres. The list an agent carries around will therefore consist of place names and trust centre names. This list is called „System Permission List" (SPL). Figure 5 explains the System Permission List as an example implementation of a trust token. The data fields have the following meaning:

trustLevel	- The level of trust given by the trust centre to the referenced agent place.
agentPlace	- The name of the agent place.
startDate	- The start date from which on the trust token is valid.
expiryDate	- The date when the trust token expires.
issuingAuthority	- The name of the issuing trust centre.
serialNumber	- A serial number for the trust token.
owner	- The owner of the agent place.

[0041] Depending on the application requirements, the trust token might have additional data fields.

[0042] Figure 6 explains the architecture of a "Agent System" in conjunction with the trust mechanism. When a agent wants to migrate to another agent place. It calls the migration operation (usually called "Go"). Execution control is then transferred to the migration manager. This module in turn calls the trust manager to ensure that the agent is allowed to migrate. The trust manager checks locally with the trust token manager whether a required trust token is already stored. It further checks with the revoke list manager to ensure the status of the trust token. If it cannot find a required token locally, it uses the trust token requester to get a trust token from the agent place, a directory service, or the trust centre. If the new trust token passes the revoke list manager test, the agent is handed to the migration handler module which performs the migration procedure.

[0043] A place can get a kind of license (a trust token) from a trust centre, if the trust centre is convinced that the place is trustworthy. To determine, if a place is trustworthy the trust centre can e.g. check the owner of the place, the software etc. One can even think of checking the everyday operation of the organisation operating the agent place. Every trust centre may have different requirements a place has to fulfil before it gets a license from it. Even a single trust centre may have different levels of such requirements resulting in different levels of licenses.

[0044] If a place fulfils the requirements of a trust centre it gets the license („Issuing of the trust token"). This kind of license is called a trust token. A trust token is a statement digitally signed by the issuing trust centre, containing the information that a specified place is trusted by the issuer and also some additional information. Trust tokens are very similar to certificates used in authentication protocols. A trust centre is very similar to a certificate authority in the same context.

[0045] Trust centres may charge places for their services.

Figure 6 explains the architecture of a sending agent system. An agent wishing to migrate to another agent place issues this request to the migration manager. The migration manager inquires a trust manager about the trustworthiness of the other agent place. The trust token requester retrieves the trust token using one of the possible ways (e.g. accessing the peer agent place, accessing the directory or accessing the issuing trust centre).

[0046] If the trust token allows the migration the mobile agent will be handed to the migration handler which executes the migration step.

[0047] If an agent now wants to migrate to a place and the place is not contained in the list of trusted places. It tries to check whether the place has a trust token from a trust centre the agent trusts. This can be done in several ways.

1. The agent may get one of the wanted trust tokens by asking the destination place.

2. He can ask the trust centres

3. He can ask other services which stores and manages trust tokens (e.g. a directory service).

[0048] If the agent can get a trust token issued by a trust centre, which is in his SPL he is allowed to migrate. Otherwise the agent is not allowed to migrate. The whole test can be done by the agent himself or by the agent place.

[0049] If the agent is allowed to migrate then there is something like a trust chain: the user trusts a trust centre, which in turn trusts the destination place.

[0050] One problem of a trust token concept is that sometimes an issued trust token has to be revoked. This is a hard problem as trust tokens are distributed without control of the trust centre. By adding an expiry date, trust tokens have a limited lifetime so that the trust centre can be sure that a trust token is not valid anymore. Furthermore, the trust centre can maintain a revoke list which contains all tokens revoked before their expiry date.

[0051] The system required for this invention consists of the trust centre, the trust token data structure and the integration of a trust manager into the agent system.

[0052] A created trust token is stored in the token manager database. This database manages all issued trust tokens together with additional information when the cite was checked, which tests have been performed, etc. The trust token is then communicated to the agent place which can add the trust token to its database of locally managed trust tokens. This database is managed by the trust token manager component of an agent place.

[0053] If a trust token has to be revoked early, the revoked trust token are stored in a database managed by the revoke manager. Note that because of the distributed nature of the system, trust token remain valid until their expire date is over. The revoke operation is just an additional mean for agent places to verify that a given trust token is still valid. As this comparison with the trust centre revoke list is not mandatory, the trust token cannot reliably be revoked early.

[0054] The trust centre can furthermore use the directory access module to store the trust token in a generally accessible directory service. Using the token requester, anybody can request an issued trust token for a specific agent place.

[0055] In an agent place, the collection of trust tokens is stored locally in a database managed by the trust to-

ken manager. If trust tokens are requested from another place, the trust token is fetched from the database and transmitted to the requester. An instance requesting a trust token can therefore use either the agent place, the directory or the trust centre itself for requesting the trust token. Other means are also possible.

Example

[0056] Agent a wants to migrate from place p1 to place p2. His SPL does not contain p2, but the trust centres tc1 and tc2. Now it must be checked if p2 owns a trust token from either tc1 or tc2. To achieve this p2 is asked to provide all trust tokens it owns. P2 owns trust tokens from the trust centres tc3, tc2 and tc4. It sends theses tokens to p1. P1 now checks whether one of the trust tokens from p2 matches a trust centre in the SPL of a. This is the case as p2 owns a trust token from tc2 and tc2 is in the SPL of a. The agent is therefore allowed to migrate.

Main advantageous differences between the invention and the state of the art

[0057] By introducing a third party we gain the advantage that the user must not know all places his agents may travel by himself without allowing the agents to visit untrusted places. Our approach therefore does not restrict the number of places as much as the original approach did.

[0058] As trust centres may be able to perform a lot of checks to a place or can make a contract with the place owner, the trust relationship between the trust centre and the place is in most cases much more justified than the trust relationship between an agent owner and the place.

[0059] The original approach had the problem that the SPL could become very large as the user specified single places. As this list has to be moved together with the agent this is not a very good thing - especially if the agent migrates over a wireless network link with limited bandwidth. According to the present approach the SPL is much shorter, as the user can specify with one entry implicitly a whole set of agent systems.

[0060] One example of an application of the present invention is a bank offering services on the base of mobile agents. Several places for mobile agents are offered. The bank is its own trust centre for these places and issues trust token for said places. A mobile agent which wants to use the services offered by the bank includes the trust token of the bank in its SPL.

[0061] The invention as shown above presents means and methods for managing trusted places in a mobile agent system. The apparatus restricts the migration operation of the mobile agents to the boundaries defined by the trusted places. The invention contains means and methods for controlling the trustworthiness of an agent place. The migration operation of a mobile

agent is restricted by a system using trusted system concepts. The system has accessed to a list of trusted systems defined by the user or trust labelling authorities. It can check whether a migration is allowed on the basis if the target node of the network is trusted or not.

[0062] The invention focuses on the way how the trustworthiness of a place is determined. In addition to the usual list of places which are allowed to be visited, a user can define a list of trust checking agencies (trust centres). These trust centres are certificating agent places for the trustworthiness. If a place is considered trustworthy, the trust checking agency is issuing a trust certificate which is stored at the agent place. Trust certificates are encoded to prevent the manipulation. Before migrating to an agent place, the mobile agent can request the trust certificates of that place and check them against its own list of trust checking agencies. If a trust certificate fits, the agent is allowed to migrate.

Claims

1. Method for managing a communication network, comprising the following steps:

   - checking (301) the trustworthiness of at least one node (102a) of the network,
   - comparing the trustworthiness of the checked node (102a) of the network with a preset trust threshold,
   - in case the trustworthiness exceeds a preset trust threshold, generating (306) a trust token for the checked node (102a) and storing (303) the trust token in the network, wherein the trust token indicates that the corresponding node (102a) has passed the trustworthiness check.

2. Method according to claim 1, characterized in that the step of checking (301) the trustworthiness of at least one node (102a) of the network is executed (306) by another node (301) of the network having a trust center function.

3. Method according to anyone of claims 1 or 2, characterized by the step of adding a predetermined expiry date to a trust token.

4. Method for managing the migration of mobile agents to nodes of a communication network, comprising the following steps:

   - checking (301) the trustworthiness of at least one node (102b) of the network,
   - in case the trustworthiness exceeds a preset trust threshold, generating (306) a trust token for the checked node (102b) and storing (303)

the trust token in the network, wherein the trust token indicates that the corresponding node (102a) has passed the trustworthiness check,
   - in advance to a migration of a mobile agent (104a) to a node of the network, verifying (108, 109) if a valid trust token is existing for the corresponding node, and
   - restricting (107) the migration of the mobile agent (104a) to nodes having a valid trust token.

5. Method according to claim 4, characterized in that the verification (108, 109) if a valid trust token is existing for the corresponding node is executed dynamically during the runtime of a mobile agent (104a).

6. Method according to claim 4 or 5, characterized in that the step of checking (301) the trustworthiness of at least one node of the network is executed by at least one third node (301) of the network having a trust center function.

7. Method according to claim 6, characterized in that a list of trust center nodes is attached to a mobile agent (104a).

8. Method according to anyone of the preceding claims, characterized in that the step of verification (108, 109) if a valid trust token is existing for the corresponding node in advance to a migration of the mobile agent (104a) is effected by

   - querying (111) the corresponding target node (102b) for valid trust token,
   - querying (112) at least one trust center node (301), and/or
   - querying (113) at least one storage unit storing valid trust token and the corresponding nodes of the network.

9. Method according to anyone of the preceding claims, characterized by the step of adding a predetermined expiry date to a trust token.

10. Software element for executing, when loaded in a computing device, a method according to anyone of the preceding claims.

11. Communication network, comprising at least one trust center node (301) for

checking the trustworthiness of at least one node of
the network, and in case the trustworthiness ex-
ceeds a preset trust threshold, the trust center node
(301) having a token creator (306) for generating a
trust token for the checked node, the trust center
node being another node of the network as each
checked node (102b),
wherein the trust token and the corresponding
checked nodes are stored in a trust manager data-
base (303) being part of or being connected to the
trust center node (301).

12. Communication network according to claim 11
    characterized by

        at least one mobile agent place (101a) manag-
        ing the migration of mobile agents in the com-
        munication network,
        the mobile agent places (101a) being designed
        to access the trust manager database (303) of
        at least one trust center node (301).

13. Communication network according to claim 12,
    characterized in that
    the mobile agent place (101a) comprises a local da-
    tabase to store information fetched from the trust
    manager database (303) of a trust center node
    (301).

14. Communication network according to anyone of
    claims 11 to 13,
    characterized in that
    a list of trust center nodes attached to a mobile
    agent (104a) managed by a mobile agent place
    (101a).

15. Communication network according to anyone of
    claims 11 to 14,
    characterized in that
    each trust token comprises an expiry date.

104a ∨ Mobile Agent

103a

System Agent

105   104a   106   103b
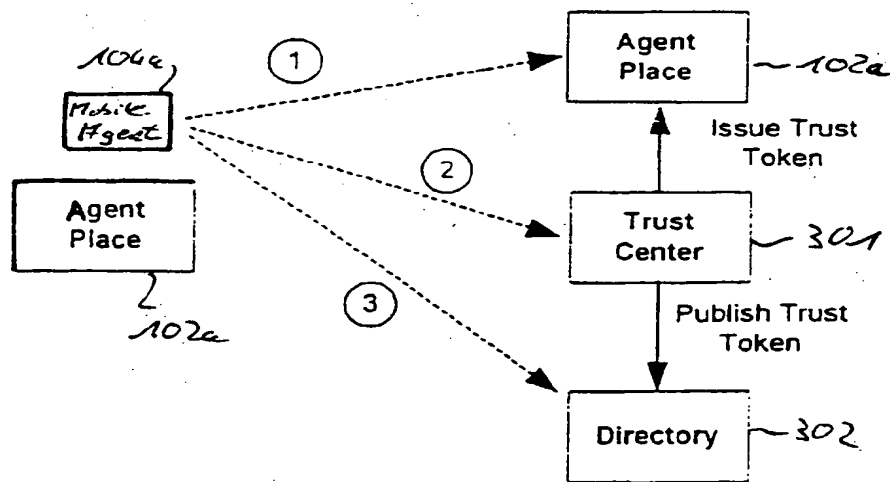
Place

102a ∨

101a ∨ Agent System

102b ∨

101b ∨ Agent System

Place

Fig 1.

Fig 2.

Fig 3.

Fig 4.

11

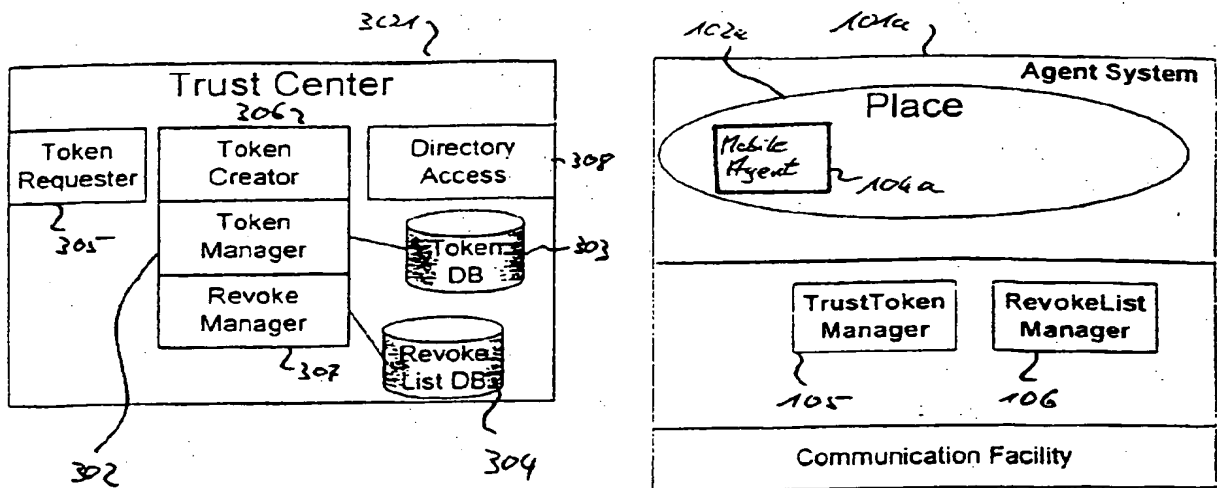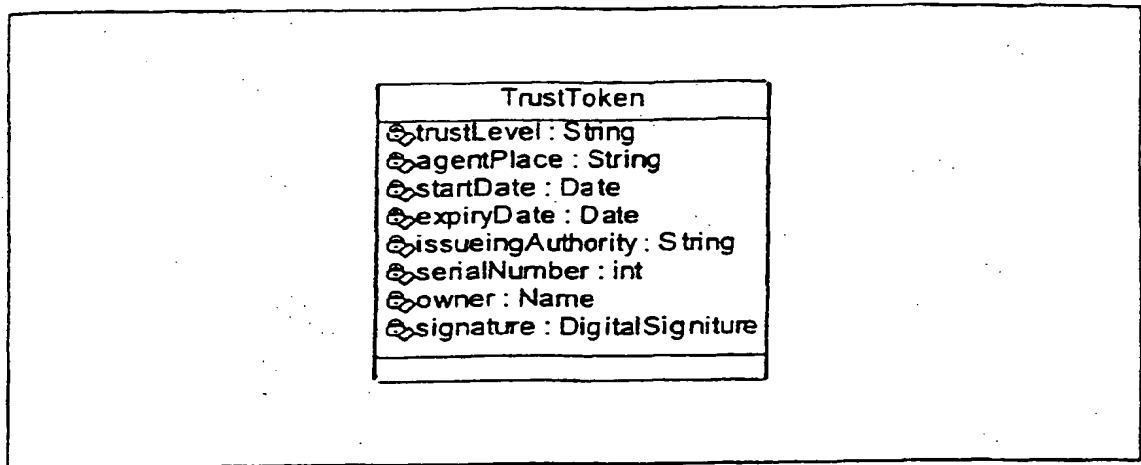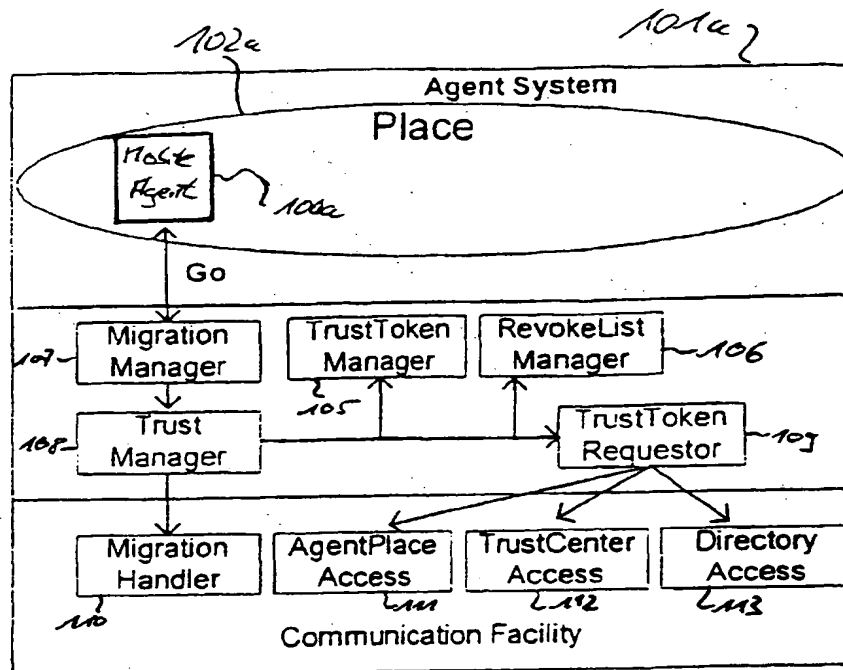| TrustToken |
|---|
| trustLevel : String |
| agentPlace : String |
| startDate : Date |
| expiryDate : Date |
| issueingAuthority : String |
| serialNumber : int |
| owner : Name |
| signature : DigitalSigniture |
| |

Fig 5.

Fig 6.

European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 99 11 2983

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
Y	ORDILLE J J: "When agents roam, who can you trust?" PROCEEDINGS. THE FIRST ANNUAL CONFERENCE ON EMERGING TECHNOLOGIES AND APPLICATIONS IN COMMUNICATIONS (CAT. NO.96TB100035), PROCEEDINGS OF COM'96. FIRST ANNUAL CONFERENCE ON EMERGING TECHNOLOGIES AND APPLICATIONS IN COMMUNICATIONS, PORTLAND, OR, USA, . pages 188-191, XP002124935 1996, Los Alamitos, CA, USA, IEEE Comput. Soc. Press, USA ISBN: 0-8186-7585-3	1-4,6,7, 9-11,14, 15	G06F9/46 H04L29/06
A	* the whole document *	5,8,12, 13	
Y	VARADHARAJAN V ET AL: "AN APPROACH TO DESIGNING SECURITY MODEL FOR MOBILE AGENT BASED SYSTEMS" IEEE GLOBAL TELECOMMUNICATIONS CONFERENCE,US,NEW YORK, NY: IEEE, page 1600-1606 XP000805177 ISBN: 0-7803-4985-7 * page 1604, left-hand column, line 36 - page 1604, left-hand column, line 36 *	1-4,6,7, 9-11,14, 15	
A	P. NIKANDER: "An architecture for authorization and delegation in distributed object-oriented agent systems" DOCTORAL DISSERTATION, 19 March 1999 (1999-03-19), pages 1-70, XP002124936 University of Helsinski, Finland * page 20, paragraph 2.1.1 - page 21 * * page 33, line 1 - line 5; figure 7 * * page 49, paragraphs 4.1,4.2 - page 56 *	1,4,11	

TECHNICAL FIELDS SEARCHED (Int.Cl.7)

G06F

-/--

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	14 December 1999	Michel, T

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

		Application Number
European Patent Office	EUROPEAN SEARCH REPORT	EP 99 11 2983

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication. where appropriate. of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	PAPAIOANNOU T ET AL: "Using mobile agents to improve the alignment between manufacturing and its IT support systems" ROBOTICS AND AUTONOMOUS SYSTEMS,NL,ELSEVIER SCIENCE PUBLISHERS, AMSTERDAM, vol. 27, no. 1-2, page 45-57 XP004164318 ISSN: 0921-8890 * page 55, right-hand column, paragraph 8.4 *	1,4,11	
A	TARDO J ET AL: "Mobile agent security and Telescript" DIGEST OF PAPERS. COMPCON '96. TECHNOLOGIES FOR THE INFORMATION SUPERHIGHWAY. FORTY-FIRST IEEE COMPUTER SOCIETY INTERNATIONAL CONFERENCE (CAT. NO.96CB35911), COMPCON '96. TECHNOLOGIES FOR THE INFORMATION SUPERHIGHWAY DIGEST OF PAPERS, SANTA CLARA, CA, pages 58-63, XP002124937 1996. Los Alamitos, CA, USA, IEEE Comput. Soc. Press, USA ISBN: 0-8186-7414-8 * the whole document *	1,4,11	TECHNICAL FIELDS SEARCHED (Int.Cl.7)

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	14 December 1999	Michel, T

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)